

Design of Resource Agents with Guaranteed Tracking Properties for Real-Time Control of Electrical Grids

Andrey Bernstein, Niek J. Bouman, and Jean-Yves Le Boudec
EPFL, Switzerland

Abstract—We focus on the problem of controlling electrical microgrids with little inertia in real time. We consider a centralized controller and a number of resources, where each resource is either a load, a generator, or a combination thereof, like a battery. The centralized controller periodically computes new power setpoints for the resources based on the estimated state of the grid and an overall objective, and subject to safety constraints. Each resource is augmented with a *resource agent* that a) implements the power-setpoint requests sent by the controller on the resource, and b) translates device-specific information about the resource into an abstract, device-independent representation and transmits this to the controller.

We focus on the resource agents (including the resources that they manage) and their impact on the overall system's behavior. Intuitively, for the system to converge to the overall objective, the resource agents should be *obedient* to the requests from the controller, in the sense that the actually implemented setpoint should be close to the requested setpoint, at least on average. This can be important especially when a controller that performs *continuous* optimization is used (because of performance constraints) to control *discrete* resources (for which the set of implementable setpoints is discrete).

In this work, we formalize this type of obedience by defining the notion of *c*-bounded accumulated-error for some constant *c*. We then demonstrate the usefulness of our notion, by presenting theoretical results (for a simplified scenario) as well as some simulation results (for a more realistic setting) that indicate that, if all resource agents in the system have *c*-bounded accumulated-error, the closed-loop system converges on average to the objective. Finally, we show how to design a resource agent that provably has *c*-bounded accumulated-error for various types of resources, such as resources with uncertainty (e.g., PV panels) and resources with a discrete set of implementable setpoints (e.g., heating systems with heaters that each can either be switched on or off).

I. INTRODUCTION

WE consider the problem of controlling a collection of electrical resources that are interconnected via an electrical grid, with respect to a (typically time-varying) objective and under certain safety constraints. These resources can be generators, loads, as well as resources that can both inject and consume power, like storage units. This problem has recently received renewed interest through the advent of renewable energy like solar power and improved battery technologies.

We focus on *real-time* control, namely on sub-second time scale. The classical approach involves a combination of both frequency and voltage regulation using *droop controllers*. With the increased penetration of stochastic resources, distributed generation and demand response, this approach shows severe

limitations in both the optimal and feasible operation of medium and low voltage networks, as well as in the aggregation of the network resources for upper-layer power systems. An alternative approach is to directly control the targeted grid by defining explicit and real-time setpoints for active/reactive power absorptions/injections defined by a solution of a specific optimization problem. Such an approach typically assumes a bi-directional communication between the *grid controller* and the different resources in the network. The communication capability enables the grid controller to have fine-grained knowledge of the system's state, which allows for a better operation of the system.

One of the major challenges in this context is to be able to efficiently control *heterogeneous resources* in real-time. The resources can have continuous or discrete nature (e.g., heating systems consisting of a finite number of heaters that each can either be switched on or off) and/or can be highly uncertain (e.g., PV panels or residential loads). Hence, a naive approach would lead to a *stochastic mixed-integer* optimization problem to be solved at the grid controller at each time step. Since the goal is real-time control, this approach is practically infeasible.

The recently introduced COMMELEC framework for real-time control of electrical grids [1], [2] inherently avoids performing stochastic mixed-integer optimization. The framework uses a hierarchical system of *software agents*, each responsible for a single resource (loads, generators and storage devices; Resource Agents - RA) or an entire subsystem (including a grid and/or a number of resources; Grid Agents - GA). Each resource agent advertises to its grid agent its internal state via a *device-independent* protocol. In particular, the protocol requires from every RA to advertise (i) a *convex* set of feasible setpoints, and (ii) an uncertainty (belief) set in setpoint implementation. This way, the grid agent can solve a *robust continuous* (rather than stochastic mixed-integer) optimization problem and send continuous setpoints to the resource agents. It is then the task of the RA to map the received setpoint to the set of actually implementable setpoints at this moment.

In this paper, we address the following two questions that arise during RA design in this context:

- (i) How should an RA advertise the flexibility and uncertainty of the resource to a GA that works with continuous setpoints?
- (ii) Given a requested setpoint from the GA (which is not necessarily implementable), which setpoint should the RA implement?

It is important to note that the RA design can have a significant impact on the performance of the overall system. Indeed, as it is shown in Section VIII, a straightforward approach in which the requested setpoint is merely projected to the set of implementable setpoints can lead to highly sub-optimal behaviour.

The key point is that the RAs should be *obedient* to the GA's requests in some sense. As the first main contribution of this paper, we propose *the boundedness of the accumulated error between requests and implemented setpoints* as a desired metric for obedience. One of the main motivations of this choice is that the average requested and implemented setpoints are the same in the long-run. This has an energy interpretation: the produced/consumed energy converges to the requested one. The latter can be useful, e.g., in *virtual power plant*-related applications. Further, we perform a theoretical analysis of the closed-loop system that includes a GA and a number of RAs, where each RA has bounded accumulated-error. The analysis is performed in a restricted scenario, under some simplifying assumptions. To complement this, we illustrate the performance of our method in simulation using a more realistic scenario.

The second main contribution of the paper is a framework for RA design that guarantees boundedness of the accumulated error, which covers a broad range of resources. Our approach is conceptually simple and is inspired by a generic error-feedback technique that is known under different names in different fields; in image processing, it is called *error diffusion* (and a specific variant is called *Floyd–Steinberg dithering* [3]), whereas in signal processing, the technique is related to sigma-delta modulation [4] and some stability properties of this technique have been analyzed in that context [5], [6].

We briefly outline the key ideas below. Consider a sequence of requested setpoints (x_k) . Assume that the resource agent cannot implement x_k precisely, but rather can implement any setpoint y_k in a given set \mathcal{I}_k (not necessarily convex, possibly discrete). Let $e_k = \sum_{i=1}^k (y_i - x_i)$ be the accumulated error. Then, roughly speaking, we answer the two questions posed above as follows. First, advertise the *convex hull* of \mathcal{I}_k as the set of feasible setpoints. Second, when asked to implement x_k , actually implement $y_k = \text{proj}_{\mathcal{I}_k}(x_k - e_{k-1})$, where $\text{proj}_{\mathcal{I}_k}(\cdot)$ is a projection operator. Under certain conditions (detailed in Section V), this method ensures boundedness of e_k .

Our method can also be interpreted from the perspective of PI control. Indeed, it can be viewed as an “I-controller” in the following sense. Observe that given the previous accumulated error e_{k-1} and request x_k , we have that y_k chosen above minimizes the absolute value of the new accumulated error $|e_k| = |e_{k-1} + (y_k - x_k)| = |y_k - (x_k - e_{k-1})|$ over all y_k in \mathcal{I}_k .

Our results hold universally, for any sequence of requested setpoints. This allows for a separation of concerns: the designer of an RA does not depend on a particular optimization algorithm applied in the GA.

The structure of the paper is as follows. In Section II, we introduce notation and briefly overview the COMMELEC framework. In Section III, we introduce the concept of the c -bounded accumulated-error and discuss its basic properties. In

Section IV, we study analytically the effect of having resource agents with bounded accumulated-error in a system where those resource agents are controlled by a grid agent. Then, we devise a method for RA design for *discrete* resource that guarantees boundedness of the accumulated error in Section V, and for *uncertain* resources in Section VI. In Section VII, we present a unified approach that allows to prove the results of Sections V and VI, as well as to design RAs for different types of resources. In Section VIII, we present a numerical evaluation of the proposed methods. Finally, we close with concluding remarks in Section IX.

II. PRELIMINARIES

In this section, we introduce some notation and also briefly review the properties of the COMMELEC framework that are used in the paper.

A. Notation

Throughout the paper, $\|\cdot\|$ denotes the ℓ_2 norm. We use \mathbb{N} and $\mathbb{N}_{>0}$ to refer to the natural numbers including and excluding zero, respectively. For arbitrary $n \in \mathbb{N}_{>0}$, we write $[n]$ for the set $\{1, \dots, n\}$. We use $\mathbf{0}$ to denote the zero vector $(0, \dots, 0)$, where the dimension should be clear from the context. Let $d \in \mathbb{N}_{>0}$. For an arbitrary set $\mathcal{U} \in \mathbb{R}^d$, $-\mathcal{U} := \{-u \mid \forall u \in \mathcal{U}\}$. For arbitrary sets $\mathcal{U}, \mathcal{V} \in \mathbb{R}^d$, $\mathcal{U} + \mathcal{V}$ represents the Minkowski sum of \mathcal{U} and \mathcal{V} , which is defined as $\mathcal{U} + \mathcal{V} := \{u + v \mid \forall u \in \mathcal{U}, \forall v \in \mathcal{V}\}$.

Fix $d \in \mathbb{N}_{>0}$. Let $\mathcal{S} \subset \mathbb{R}^d$ be an arbitrary non-empty closed set. Any mapping $\text{proj}_{\mathcal{S}} : \mathbb{R}^d \rightarrow \mathcal{S}$ that satisfies

$$\text{proj}_{\mathcal{S}}(x) = s, \quad \text{where } s \in \arg \min_{\rho \in \mathcal{S}} \|\rho - x\|.$$

is called a *projection operator* onto \mathcal{S} .

A *power setpoint* is a tuple $u = (P, Q) \in \mathbb{R}^2$, where P denotes real power and Q denotes reactive power. Let $\mathcal{B}(\mathbb{R}^2)$ denote the collection of all *bounded* non-empty closed subsets of \mathbb{R}^2 , and let $\mathcal{C}(\mathbb{R}^2)$ denote the collection of all *convex* sets in $\mathcal{B}(\mathbb{R}^2)$.

For vectors $v \in \mathbb{R}^2$, as well as for sets $\mathcal{V} \subseteq \mathbb{R}^2$, we use superscripts P and Q to refer to the first and second coordinate, respectively. I.e., $v = (v^P, v^Q)$ and $\mathcal{V} = \mathcal{V}^P \times \mathcal{V}^Q$ hold.

For any compact set \mathcal{V} , we define the *diameter* of \mathcal{V} as

$$\text{diam } \mathcal{V} := \max\{\|v - w\| : v, w \in \mathcal{V}\}.$$

For any matrix M , we write $M \succeq 0$ to state that it is *positive semi-definite*, and $M \succ 0$ to state that it is *positive definite*. Also, we write $\|M\|$ to denote its *induced ℓ_2 matrix norm*.

B. The COMMELEC Framework

In the COMMELEC framework [1], the RAs periodically advertise to the GA the set of (power) setpoints that they can currently implement, as well as, for each setpoint in this set, its “cost” (which exposes the individual objective of the resource) and the accuracy with which the setpoint can be implemented. The GA collects those advertisements and uses them to periodically compute new setpoints, which are

then sent back to the resources as requests, i.e., the RAs are supposed to implement these setpoints. (Note that the setpoint-computation might depend on auxiliary inputs, such as the state of the electrical grid.) Below, we present the advertised elements in more detail.

Definition 1. A *PQ Profile* is a convex set $\mathcal{A} \in \mathcal{C}(\mathbb{R}^2)$.

A *PQ* profile of a RA represents the collection of power setpoints that the RA is able to implement. As explained in the Introduction, the convexity requirement on the *PQ* profile can be viewed as a limitation that originates from the control algorithm that is currently used in the grid agent.¹

Definition 2. A *Belief Function* is a set-valued function

$$BF : \mathcal{A} \rightarrow \mathcal{B}(\mathbb{R}^2).$$

For every setpoint $u \in \mathcal{A}$, the belief function represents the uncertainty in this setpoint implementation: when the RA is requested to implement a setpoint $u \in \mathcal{A}$, the RA states that the actually implemented setpoint (which could depend on external factors, for example on the weather in case of a PV) lies in the set $BF(u)$.

Definition 3. A *Virtual Cost Function* is a continuously differentiable function $CF : \mathcal{A} \rightarrow \mathbb{R}$.

The virtual cost function represents the RA's aversion (corresponding to a high cost) or preference (respectively, low cost) towards a given setpoint. For example, a battery agent whose battery is fully charged will assign high cost to setpoints that correspond to further charging the battery. The adjective *virtual* makes clear that we do not mean a monetary value, however, from now on we will omit this adjective and simply write “cost function”.

As the objects defined above may change at every COM-MELEC cycle, we denote the cycle (or step) index by using subscript $k = 1, 2, \dots$ where needed, e.g., CF_k , BF_k , etc.

At every time step $k \in \mathbb{N}_{>0}$, a resource agent receives from its grid agent a *request* to implement a setpoint $u_k^{\text{req}} = (P_k^{\text{req}}, Q_k^{\text{req}}) \in \mathcal{A}_k$. We denote the setpoint that the resource *actually implements* by $u_k^{\text{imp}} = (P_k^{\text{imp}}, Q_k^{\text{imp}}) \in \mathcal{I}_k$. Here, $\mathcal{I}_k \subseteq \mathbb{R}^2$ denotes the set of *implementable setpoints* at time step k , namely the set describing the *feasibility constraints* of the resource.

Note that the implemented setpoint u_k^{imp} need not be equal to the request u_k^{req} ; there is typically some error between them. For example, there might be setpoints in the *PQ* profile that do not correspond to implementable setpoints (we will see examples of this in Section V). Also, the setpoint that is actually implemented might depend on uncertain external factors (like the solar irradiance, in case of a PV).

C. Policy for a Resource Agent

It is convenient to formulate the decision process at a resource agent as a *repeated game* between the RA and the

“environment”, which includes the effects of the decisions of the GA, Nature, and other external factors. To that end, let

$$\mathcal{H}_k = \left\{ (\mathcal{A}_i, CF_i, BF_i, u_i^{\text{imp}}), (u_i^{\text{req}}, \mathcal{I}_i) \right\}_{i=1}^k \quad (1)$$

denote the *history* of the decision process up to (and including) time step k . Here, the first tuple represents an action of the RA, whereas the second tuple represents an action of the environment. A general *policy* of the RA is the collection $\pi = \{\pi_k\}_{k=1}^\infty$ of probability distributions $\pi_k = (\pi_k^{\mathcal{A}}, \pi_k^{CF}, \pi_k^{BF}, \pi_k^{\text{imp}})$, such that

$$\begin{aligned} \mathcal{A}_k &\sim \pi_k^{\mathcal{A}}(\cdot | \mathcal{H}_{k-1}), & CF_k &\sim \pi_k^{CF}(\cdot | \mathcal{H}_{k-1}), \\ BF_k &\sim \pi_k^{BF}(\cdot | \mathcal{H}_{k-1}), & u_k^{\text{imp}} &\sim \pi_k^{\text{imp}}(\cdot | \mathcal{H}_k). \end{aligned}$$

Note that, in general, the first three elements are chosen according to the history up to, *but not including*, time step k , while the implemented setpoint is chosen according to the history *including* time step k .

In this context, we make the following important distinction between *deterministic* and *uncertain* resources.

Definition 4. We say that a resource is *deterministic* if, for any time step $k \in \mathbb{N}_{>0}$, there exists a deterministic function G_k (that is known to the RA) such that $\mathcal{I}_k = G_k(\mathcal{H}_{k-1})$. Otherwise, we say that the resource is *uncertain*.

In particular, an agent for a deterministic resource knows the set of implementable setpoints \mathcal{I}_k at the moment of advertising the *PQ* profile \mathcal{A}_k , hence its policy may depend on it explicitly. On the other hand, an agent for an uncertain resource will have to *predict* \mathcal{I}_k in order to advertise \mathcal{A}_k .

III. BOUNDED ACCUMULATED-ERROR

In this section, we introduce the concept of *c*-bounded accumulated-error and we moreover propose that bounded accumulated-error should be a required property of every resource agent design. The motivation for our proposal is two-fold. First, if an RA has bounded accumulated-error, the average implemented setpoint converges to the average requested setpoint (see Proposition 1 below). Second, a resource that introduces a significant error when implementing a setpoint can lead to divergence and/or sub-optimal behaviour of the overall system. We show that, under the bounded accumulated-error property, the effect on the setpoints computed by the grid agent is bounded and vanishes on average. We prove this analytically in a restricted scenario (Section IV) and in simulation for a more realistic scenario (Section VIII).

Let

$$e_k = (e_k^P, e_k^Q) := \sum_{i=1}^k (u_i^{\text{imp}} - u_i^{\text{req}}), \quad k \in \mathbb{N} \quad (2)$$

denote the *accumulated-error vector* at time k . We define the following performance metric in terms of e_k .

Definition 5. Let $c \in \mathbb{R}$, $0 \leq c < \infty$, be given. We say that a resource agent has *c*-bounded accumulated-error if

$$\|e_k\| \leq c, \quad k \geq 1.$$

¹This particular control algorithm computes orthogonal projections onto the *PQ* profile, which are well-defined only if that set is convex.

Finally, when we say that a resource agent has *bounded accumulated-error* (without mentioning a constant), we mean that there exists some constant $c \geq 0$ such that the resource agent has c -bounded accumulated-error.

A. Vanishing Average Tracking-Error

It is easy to show that if a resource agent has bounded accumulated-error, the average implemented setpoint converges to the average requested setpoint. In fact, we next prove a stronger result that gives the rate of convergence when the average is taken over a window of a given size. To this end, for a given sequence $x = (x_k)$, any $k = 1, 2, \dots$, and $m \in \{0, \dots, k-1\}$, let

$$\mu_{[k-m:k]}(x) := \frac{1}{m+1} \sum_{i=0}^m x_{k-i}, \quad (3)$$

denote the average of x over the last $m+1$ elements. The *average tracking error* between two sequences $x = (x_k)$ and $y = (y_k)$ is then defined as

$$\varepsilon_{[k-m:k]}^{\text{avg}}(x, y) := \left\| \mu_{[k-m:k]}(x) - \mu_{[k-m:k]}(y) \right\|. \quad (4)$$

Proposition 1. *Let R be a resource agent having c -bounded accumulated-error. Then, for any $k = 1, 2, \dots$ and $m \in \{0, \dots, k-1\}$, the average tracking error of R satisfies*

$$\varepsilon_{[k-m:k]}^{\text{avg}}(u^{\text{req}}, u^{\text{imp}}) \leq \begin{cases} (m+1)^{-1}c & \text{if } m = k-1, \\ 2(m+1)^{-1}c & \text{otherwise.} \end{cases}$$

In particular, for the case $m = 0$, we obtain an upper bound on the instantaneous error

$$\|u_k^{\text{imp}} - u_k^{\text{req}}\| \leq 2c.$$

Proof:

$$\begin{aligned} \varepsilon_{[k-m:k]}^{\text{avg}}(u^{\text{imp}}, u^{\text{req}}) &= \left\| \mu_{[k-m:k]}(u^{\text{req}}) - \mu_{[k-m:k]}(u^{\text{imp}}) \right\| \\ &= \frac{1}{m+1} \left\| \sum_{i=0}^m (u_{k-i}^{\text{imp}} - u_{k-i}^{\text{req}}) \right\| \\ &= \frac{1}{m+1} \left\| \sum_{i=1}^k (u_i^{\text{imp}} - u_i^{\text{req}}) - \sum_{j=1}^{k-m-1} (u_j^{\text{imp}} - u_j^{\text{req}}) \right\| \\ &= \frac{1}{m+1} \|e_k - e_{k-m-1}\| \leq \frac{1}{m+1} (\|e_k\| + \|e_{k-m-1}\|) \end{aligned}$$

We obtain the statement by using the bound on the error from Definition 5 twice. For the case $m = k-1$, we use the fact that $e_0 = (0, 0)$ (by definition of e_k) which gives us the improved bound. ■

B. Example: Resource with Delay

Here, we would like to give some simple examples of resources that have bounded accumulated-error; in Sections V and VI, we study more elaborate examples. A trivial example is an *ideal* device, for which $u_k^{\text{imp}} = u_k^{\text{req}}$ for every k , which immediately implies that $e_k = e_0$ for every k .

Another example of a resource with bounded accumulated error by construction is a device that has a *delay* of $\tau \in \mathbb{N}_{>0}$ time steps when implementing a setpoint. For simplicity,

consider a resource with a fixed *convex* set of implementable setpoints $\mathcal{I} \subseteq \mathbb{R}^2$. Namely, the set \mathcal{I} represents the feasibility constraint of the resource. When asked at time step k to implement a setpoint $u_k^{\text{req}} \in \mathcal{I}$, $u_k^{\text{req}} \neq u_{k-1}^{\text{imp}}$, it will only implement it after τ timesteps; $u_{k+\tau}^{\text{imp}} = u_k^{\text{req}}$. During the “transition period”, the implemented setpoint remains unchanged, i.e., $u_{k+i}^{\text{imp}} = u_{k-1}^{\text{imp}}$ for all $i \in \{0, \dots, \tau-1\}$. At timesteps $k+1$ up to and including $k+\tau-1$, the PQ profile of the resource is a singleton set that corresponds to the *implemented* setpoint at the previous timestep (i.e., equal to $\{u_{k-1}^{\text{imp}}\}$); while at timestep $k+\tau$, the PQ profile is set back to \mathcal{I} . Note that no error is accumulated at timesteps where the PQ profile is a singleton set. The idea is illustrated in Figure 1.




Fig. 1. Illustration of operation of a real-power resource with delay $\tau = 3$.

We show below that this resource agent has c -bounded accumulated error with $c = \text{diam } \mathcal{I}$. Fix $k \in \mathbb{N}_{>0}$. Let $\{i_\ell\}_{\ell=1}^L$, $1 < i_1 < i_2 < \dots < i_L \leq k$, be the set of time indices such that the request differs from the implemented setpoint, namely

$$\begin{aligned} u_i^{\text{req}} &= u_i^{\text{imp}}, & \text{for every } i \in [k] \setminus \{i_\ell\}_{\ell=1}^L \\ u_{i_\ell}^{\text{req}} &\neq u_{i_\ell}^{\text{imp}}, & \text{for every } \ell \in [L]. \end{aligned}$$

Observe that because of the delay, we have in particular that $u_{i_\ell}^{\text{imp}} = u_{i_\ell-1}^{\text{imp}}$. When the request differs from the implemented setpoint, the PQ profile is “locked” on the singleton $\{u_{i_\ell-1}^{\text{imp}}\}$ during the transition period, and as a result $u_{i_\ell+i}^{\text{req}} = u_{i_\ell+i}^{\text{imp}}$, $i = 1, \dots, \tau$. At the end of the transition period, the implemented setpoint equals to the request right before the transition started, namely $u_{i_\ell+\tau}^{\text{imp}} = u_{i_\ell}^{\text{req}}$. Also, by the definition of the time indices $\{i_\ell\}_{\ell=1}^L$, it holds that $u_{i_\ell+\tau}^{\text{imp}} = u_{i_{\ell+1}}^{\text{imp}}$. It is then clear that

$$u_{i_{\ell+1}}^{\text{imp}} = u_{i_\ell}^{\text{req}}, \quad \ell \in [L]. \quad (5)$$

Thus, the accumulated error in Eqn. (2) becomes

$$e_k = \sum_{i=1}^k (u_i^{\text{imp}} - u_i^{\text{req}}) = \sum_{\ell=1}^L (u_{i_\ell}^{\text{imp}} - u_{i_\ell}^{\text{req}}) = u_{i_1}^{\text{imp}} - u_{i_L}^{\text{req}},$$

where the last equality follows by (5). Hence,

$$\|e_k\| \leq \|u_{i_1}^{\text{imp}} - u_{i_L}^{\text{req}}\| \leq \text{diam } \mathcal{I}.$$

IV. CONVERGENCE ANALYSIS OF THE CLOSED-LOOP SYSTEM

In this section, we study analytically the effect of having resource agents with bounded accumulated-error in a system where those resource agents are controlled by a grid agent.

For a restricted scenario (a set of conditions that we state precisely further below), we show that the control algorithm used in the grid agent converges to the optimum (minimum) value of the objective function.




Fig. 2. Setup involving a grid agent (GA), resource agents (RA_i), the physical resources themselves (PR_i), the electrical grid including phasor measurement units (PMUs) and the state estimator.

A. Control Algorithm: Projected Gradient Descent

Consider a grid agent that has n follower agents and possibly a leader agent. Let $x_k = (u_{k,1}^{\text{req}}, \dots, u_{k,n}^{\text{req}}) \in \mathbb{R}^{2n}$ denote the vector of requested setpoints computed by the grid agent at time step k . Similarly, let $y_k = (u_{k,1}^{\text{imp}}, \dots, u_{k,n}^{\text{imp}}) \in \mathbb{R}^{2n}$ denote the vector of actually implemented setpoints (implemented by the resources) at time step k . Let $\hat{y}_k = ((P_1, Q_1), \dots, (P_n, Q_n)) \in \mathbb{R}^{2n}$ denote the vector of estimated bus powers of the n buses in the grid to which the followers are connected, where we suppose that each resource is connected to a distinct bus, and that no other loads or generators are connected to those buses. Figure 2 shows a block diagram of this setup.

As described in [1], the grid agent computes setpoints using a continuous gradient steering algorithm

$$x_{k+1} = \text{proj}_{\mathcal{U}}(\hat{y}_k - \alpha \nabla J_k(\hat{y}_k)), \quad (6)$$

where \mathcal{U} is the set of *admissible* setpoints as defined in [1], J_k is the current objective function that includes the weighted sum of the followers cost functions as well as the penalty terms related to the grid quality of supply and deviation from the request from the leader (the parent grid agent), and α is the gradient step size. Observe that since the steering is performed from the estimated grid state \hat{y}_k , without further knowledge on the relation between x_k and \hat{y}_k , no guarantees can be provided for this algorithm. However, if the resource agents are obedient in the sense of Definition 5, we can show that under certain conditions the grid's state converges to the optimal state.

B. A Restricted Scenario

We perform the analysis of (6) under the following set of conditions.

Hypothesis 1. Consider the setting where:

- (i) The state is estimated perfectly, i.e., $\hat{y}_k = y_k$, hence the gradient steering algorithm simplifies to:

$$x_{k+1} = \text{proj}_{\mathcal{U}}(y_k - \alpha \nabla J_k(y_k)).$$

- (ii) The projection in (6) is not “active”, i.e.,

$$y_k - \alpha \nabla J_k(y_k) \in \mathcal{U} \quad \text{for all } k.$$

- (iii) The objective function J_k is fixed, namely $J_k \equiv J$.
- (iv) The objective function J is quadratic and convex, namely

$$J(x) = x^T \Gamma x + \gamma x + a$$

for some positive definite and symmetric matrix $\Gamma \succ 0 \in \mathbb{R}^{2n \times 2n}$, vector $\gamma \in \mathbb{R}^{2n}$, and $a \in \mathbb{R}$.

Condition (i) of Hypothesis 1 expresses that the state of the grid is solely determined by the setpoints implemented by the followers, and that the grid agent knows this state exactly. We leave the analysis of the more realistic “noisy case” for further research. Condition (ii) corresponds to an evolution of the system in which all operating points of the grid lie well inside the feasibility region defined by quality-of-supply constraints and the PQ profiles of the different resources. We cannot argue that this condition is always satisfied in practice — on the contrary, a main feature of the COMMELEC framework is that it can cope with grids “under stress”. What we can argue, however, is that in such stressed-grid situations, *safety*, rather than optimal performance (with respect to J_k), is the primary concern. Safety of the grid is guaranteed when the vector of setpoints lies inside the admissible set \mathcal{U} , as discussed in [1]. If, on the other hand, the grid is not stressed, then we are primarily interested in achieving (close to) optimal performance, and for this case Condition (ii) will be satisfied.

Condition (iii) merely requires that the rate of change of the advertisement messages is much slower than COMMELEC’s temporal resolution. Finally, Condition (iv) can be satisfied by requiring the resource agents to send only quadratic cost functions, and by linearizing the penalty term that is related to the grid QoS. In particular, this condition is satisfied in the setup considered in [1].

Regardless of whether Hypothesis 1 is valid in realistic scenarios, our result may be viewed as a potential stepping stone to a more general result, and as a sanity check: suppose instead that we had found an impossibility result in this simplified scenario, then this would already have ruled out the possibility of obtaining a more general convergence result.

C. Convergence Result

For the restricted scenario described above, we prove the following result.

Theorem 1. Consider a system consisting of:

- a grid agent, with control law (6) under the conditions stated in Hypothesis 1 and with step size $\alpha \leq 1/\rho(\Gamma)$, where $\rho(\Gamma)$ is the spectral radius of Γ ;
- $n \in \mathbb{N}_{>0}$ follower resource agents, where for every $i \in [n]$, resource agent i has c_i -bounded accumulated-error. Let $c := [c_1 \ \dots \ c_n] \in \mathbb{R}^n$.

Then, it holds that

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k y_k = x^*, \quad \text{where } x^* := \arg \min_x J(x) = -\Gamma^{-1} \gamma.$$

Moreover, the rate of convergence is given by

$$\left\| \frac{1}{k} \sum_{i=1}^k y_k - x^* \right\| \leq \frac{(1-r^k)}{k(1-r)} (\|\psi\| + \|\Gamma^{-1}\| \|\gamma\| r^2 + \|c\|)$$

where $r := \rho(I - \alpha \Gamma) < 1$ and $\psi := x_1 + 2\alpha \gamma - \alpha^2 \Gamma \gamma$.

We prove this theorem using the following two lemmas.

Lemma 1. Under the conditions from Hypothesis 1, the implemented setpoint y_k satisfies the following recursion

$$y_{k+1} = (I - \alpha \Gamma) y_k - \alpha \gamma + \epsilon_{k+1}, \quad k \in \mathbb{N}, k \geq 1, \quad (7)$$

where for every $k \in \mathbb{N}$

$$\epsilon_k := \begin{cases} y_k - x_k & \text{if } k \in \mathbb{N}, k \geq 1 \\ \mathbf{0} & \text{if } k = 0 \end{cases}$$

is the instantaneous error. Moreover, y_k is given explicitly by

$$y_k = (I - \alpha\Gamma)^{k-1}\psi + \sum_{i=0}^k (I - \alpha\Gamma)^{k-i}\epsilon_i - \sum_{j=0}^k (I - \alpha\Gamma)^j(\alpha\gamma) \quad (8)$$

$$= (I - \alpha\Gamma)^{k-1}\psi + \sum_{i=0}^k (I - \alpha\Gamma)^{k-i}\epsilon_i - \Gamma^{-1}(I - (I - \alpha\Gamma)^{k+1})\gamma, \quad (9)$$

for every $k \in \mathbb{N}, k \geq 1$, where $\psi := x_1 + 2\alpha\gamma - \alpha^2\Gamma\gamma$ is a constant vector that depends on the initial condition, i.e., the first request x_1 .

Proof: First observe that under Hypothesis 1, the algorithm (6) becomes

$$x_{k+1} = y_k - \alpha(\Gamma y_k + \gamma) = (I - \alpha\Gamma)y_k - \alpha\gamma.$$

Also, by definition, $x_{k+1} = y_{k+1} - \epsilon_{k+1}$. Hence, we obtain recursion (7). Next, we show that (8) satisfies (7). Indeed, plugging (8) in the right-hand-side of (7) yields

$$\begin{aligned} & (I - \alpha\Gamma)^k\psi + \sum_{i=0}^k (I - \alpha\Gamma)^{k-i+1}\epsilon_i - \sum_{j=0}^k (I - \alpha\Gamma)^{j+1}(\alpha\gamma) \\ & \quad - \alpha\gamma + \epsilon_{k+1} \\ &= (I - \alpha\Gamma)^k\psi + \sum_{i=0}^{k+1} (I - \alpha\Gamma)^{(k+1)-i}\epsilon_i - \sum_{j=0}^{k+1} (I - \alpha\Gamma)^j(\alpha\gamma) \\ &= y_{k+1} \end{aligned}$$

as required. We obtain the expression for ψ by evaluating our explicit expression for y_k at $k = 1$, i.e.,

$$y_1 = \psi + \epsilon_1 + (I - \alpha\Gamma)\epsilon_0 - \alpha\gamma - (I - \alpha\Gamma)(\alpha\gamma),$$

and then solving for ψ and using that $\epsilon_0 = \mathbf{0}$ and that $y_1 - \epsilon_1 = x_1$.

Finally, (9) follows by using the formula for geometric series of matrices. Namely, for any square matrix A we have that

$$\sum_{i=0}^k A^i = (I - A)^{-1}(I - A^{k+1})$$

provided that $I - A$ is invertible. In our case, $A = I - \alpha\Gamma$, and $I - A = \alpha\Gamma$ is invertible by Hypothesis 1 (iv). ■

Lemma 2. We have that $(I - \alpha\Gamma) \succeq 0$ and $\rho(I - \alpha\Gamma) < 1$ if and only if $\Gamma \succ 0$ and $\rho(\Gamma) \leq 1/\alpha$.

Proof: Let λ_Γ be an eigenvalue of Γ . We have that

$$\begin{aligned} 0 &< \lambda_\Gamma \leq \frac{1}{\alpha} \\ \Leftrightarrow 0 &> -\alpha\lambda_\Gamma \geq -1 \\ \Leftrightarrow 1 &> 1 - \alpha\lambda_\Gamma \geq 0. \end{aligned}$$

However, $1 - \alpha\lambda_\Gamma$ is an eigenvalue of $I - \alpha\Gamma$ if and only if λ_Γ is an eigenvalue of Γ , which completes the proof. ■

Proof of Theorem 1: We use Lemma 1, and take the average of y_1, \dots, y_k :

$$\begin{aligned} \frac{1}{k} \sum_{i=1}^k y_i &= \frac{1}{k} \sum_{i=1}^k \left[(I - \alpha\Gamma)^{i-1}c_1 + \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i \right. \\ &\quad \left. - \Gamma^{-1}(I - (I - \alpha\Gamma)^{i+1})\gamma \right] \\ &= -\Gamma^{-1}\gamma + \left(\frac{1}{k} \sum_{i=1}^k (I - \alpha\Gamma)^{i-1} \right) c_1 \\ &\quad + \Gamma^{-1} \left(\frac{1}{k} \sum_{i=1}^k (I - \alpha\Gamma)^{i+1} \right) \gamma \\ &\quad + \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i. \end{aligned}$$

Note that $x^* = -\Gamma^{-1}\gamma$ and set $r := \rho(I - \alpha\Gamma)$. We have that

$$\begin{aligned} & \left\| \frac{1}{k} \sum_{i=1}^k y_i - x^* \right\| \\ & \leq \frac{\|c_1\|}{k} \sum_{i=1}^k \|(I - \alpha\Gamma)^{i-1}\| + \frac{\|\Gamma^{-1}\| \|\gamma\|}{k} \sum_{i=1}^k \|(I - \alpha\Gamma)^{i+1}\| \\ & \quad + \left\| \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i \right\| \\ &= \frac{\|c_1\|}{k} \sum_{i=1}^k \rho((I - \alpha\Gamma)^{i-1}) + \frac{\|\Gamma^{-1}\| \|\gamma\|}{k} \sum_{i=1}^k \rho((I - \alpha\Gamma)^{i+1}) \\ & \quad + \left\| \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i \right\| \\ &= \frac{\|c_1\|}{k} \sum_{i=1}^k r^{i-1} + \frac{\|\Gamma^{-1}\| \|\gamma\|}{k} \sum_{i=1}^k r^{i+1} \\ & \quad + \left\| \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i \right\| \\ &= \frac{\|c_1\| (1 - r^k)}{k(1 - r)} + \frac{\|\Gamma^{-1}\| \|\gamma\| r^2 (1 - r^k)}{k(1 - r)} \\ & \quad + \left\| \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j}\epsilon_i \right\|. \end{aligned} \quad (10)$$

In the above sequence, the first equality follows by the fact that for any symmetric matrix A , $\|A\|_2 = \rho(A)$; and the second equality holds by the fact that $\rho(A^i) = (\rho(A))^i$. (The latter holds for an arbitrary A .) Observe that by the hypothesis of this theorem and Lemma 2, $r < 1$. Thus, the first two terms

in (10) converge to zero. As for the third term, we have

$$\begin{aligned}
& \left\| \frac{1}{k} \sum_{i=1}^k \sum_{j=0}^i (I - \alpha\Gamma)^{i-j} \epsilon_i \right\| = \\
& \frac{1}{k} \left\| [(I - \alpha\Gamma)\epsilon_0 + \epsilon_1] + [(I - \alpha\Gamma)^2\epsilon_0 + (I - \alpha\Gamma)\epsilon_1 + \epsilon] + \dots + [(I - \alpha\Gamma)^k\epsilon_0 + \dots + \epsilon_k] \right\| \\
& = \frac{1}{k} \left\| \sum_{i=1}^k \epsilon_i + (I - \alpha\Gamma) \sum_{i=0}^{k-1} \epsilon_i + (I - \alpha\Gamma)^2 \sum_{i=0}^{k-2} \epsilon_i \right. \\
& \quad \left. + \dots + (I - \alpha\Gamma)^k \epsilon_0 \right\| \\
& \leq \frac{1}{k} \left(\left\| \sum_{i=1}^k \epsilon_i \right\| + r \left\| \sum_{i=0}^{k-1} \epsilon_i \right\| + r^2 \left\| \sum_{i=0}^{k-2} \epsilon_i \right\| + \dots + r^k \|\epsilon_0\| \right) \\
& \leq \frac{\|c\|}{k} \sum_{i=0}^{k-1} r^i = \frac{\|c\| (1 - r^k)}{k(1 - r)},
\end{aligned}$$

where the second inequality follows by the premise that all resource agents have bounded accumulated-error, and the fact that $\epsilon_0 = \mathbf{0}$. Combining this with (10) completes the proof of the Theorem. ■

V. BOUNDED ACCUMULATED-ERROR: DISCRETE (REAL-POWER) RESOURCES

In this section, we show how to achieve boundedness of the accumulated error for resources that can only implement power setpoints from a *discrete* set. As an application, we consider a heating system consisting of a finite number of heaters that each can either be switched on or off (see Section V-C below). We focus here on *deterministic* resources in the sense of Definition 4. We defer the treatment of uncertain resources to Section VI.

A. A General Construction

To keep the exposition simple, and because it suffices for our concrete examples presented in Section V-C, we restrict here to a scenario with a resource that only produces/consumes *real power*.

Definition 6. For any finite non-empty set $\mathcal{S} \subset \mathbb{R}$, whose elements we label as $s_1 < s_2 < \dots < s_{|\mathcal{S}|}$, we define the maximum stepsize of \mathcal{S} as

$$\Delta_{\mathcal{S}} := \begin{cases} 0 & \text{if } |\mathcal{S}| = 1 \\ \max_{i \in [|\mathcal{S}|-1]} s_{i+1} - s_i & \text{if } |\mathcal{S}| > 1. \end{cases}$$

Definition 7. For every finite non-empty collection of sets $\mathbb{S} := \{\mathcal{S}_i\}_{i \in [n]}$ where $n \in \mathbb{N}$ and where $\mathcal{S}_i \subset \mathbb{R}$ is a finite non-empty set for all $i \in [n]$, we define the maximum step size of \mathbb{S} as

$$\Delta_{\mathbb{S}} := \max_{i \in [n]} \Delta_{\mathcal{S}_i}.$$

The following theorem states our main result for discrete real-power resources. Its proof is deferred to Section VII.

Theorem 2. Let \mathbb{S} be a finite non-empty collection of sets $\mathbb{S} := \{\mathcal{S}_i\}_{i \in [n]}$ where $n \in \mathbb{N}_{>0}$ and $\mathcal{S}_i \subset \mathbb{R}$ is a finite non-empty set for all $i \in [n]$. Let $\Delta_{\mathbb{S}}$ denote the maximum stepsize of \mathbb{S} . Let $(i_k)_{k \in \mathbb{N}_{>0}}$ be any sequence with $i_k \in [n]$, and let $P_k^{req} \in \text{conv}(\mathcal{S}_{i_k})$ for every $k \in \mathbb{N}_{>0}$. Let $e_k^P \in \mathbb{R}$ be the accumulated error as defined in Eqn. (2) for every $k \in \mathbb{N}$. Then, if a resource agent implements

$$P_k^{imp} = \text{proj}_{\mathcal{S}_{i_k}}(P_k^{req} - e_{k-1}^P) \quad (11)$$

for every $k \in \mathbb{N}_{>0}$, it has $\frac{1}{2}\Delta_{\mathbb{S}}$ -bounded accumulated-error.

For a discrete-resource agent that implements setpoints according to Theorem 2, it follows immediately from Proposition 1 that

$$|P_k^{imp} - P_k^{req}| \leq \Delta_{\mathbb{S}}.$$

For a special case that we define below, we have a slightly stronger result that proves optimality, in some sense, of our construction.

B. Accuracy of Setpoint Implementation

Definition 8. Let $\mathcal{S} \subset \mathbb{R}$ be as in Definition 6. We say that \mathcal{S} is *uniform* if $\Delta_{\mathcal{S}} = s_{i+1} - s_i$ for all $i \in [|\mathcal{S}| - 1]$. Similarly, we say that a collection $\mathbb{S} := \{\mathcal{S}_i\}_{i \in [n]}$ is *uniform* if \mathcal{S}_i is uniform and $\Delta_{\mathcal{S}_i} = \Delta_{\mathbb{S}}$ for all $i \in [n]$.

In addition, we require the following definition of a special projection operator.

Definition 9. Fix $d \in \mathbb{N}_{>0}$. For any non-empty closed set $\mathcal{S} \subset \mathbb{R}^d$ we define

$$\text{proj}_{\mathcal{S}}^y(x) := \arg \min_{\rho \in \mathcal{T}(\mathcal{S}, x)} \|\rho - y\| \quad x, y \in \mathbb{R}^d$$

where

$$\mathcal{T}(\mathcal{S}, x) := \arg \min \{\|\sigma - x\| : \sigma \in \mathcal{S}\}$$

Informally speaking, this projection operator is such that, in case the cardinality of $\mathcal{T}(\mathcal{S}, x)$ is strictly larger than one (which happens when the argument x has exactly the same distance to multiple points in \mathcal{S}), it chooses the point in $\mathcal{T}(\mathcal{S}, x)$ that is closest to some given point y .

Theorem 3. Consider the setting of Theorem 2. Assume in addition that: (i) \mathbb{S} is uniform and $\Delta_{\mathbb{S}} > 0$, and (ii) a resource agent implements

$$P_k^{imp} = \text{proj}_{\mathcal{S}_{i_k}}^{P_k^{req}}(P_k^{req} - e_{k-1}^P)$$

for every $k \in \mathbb{N}_{>0}$, where the projection operator is given by Definition 9. Then, the RA has $\frac{1}{2}\Delta_{\mathbb{S}}$ -bounded accumulated-error, and, additionally, the following properties hold:

- (i) $|P_k^{imp} - P_k^{req}| < \Delta_{\mathbb{S}}$ holds for every k (with strict inequality), and this is the best possible bound in the following sense: For any algorithm that achieves the bounded accumulated-error property and any $\varepsilon > 0$, there exists a sequence (P_k^{req}) and a sequence (i_k) such that $|P_{k_0}^{imp} - P_{k_0}^{req}| \geq \Delta_{\mathbb{S}} - \varepsilon$ for some k_0 .

- (ii) If $P_{k_o}^{\text{req}} \in \mathcal{S}_{i_{k_o}}$ holds for some k_o (in words: if the request is implementable itself), then $P_{k_o}^{\text{imp}} = P_{k_o}^{\text{req}}$, regardless of the value of the accumulated error e_{k_o-1} .

Proof: Observe that by Theorem 2, $\|e_{k-1}\| \leq \frac{1}{2}\Delta_S$ for every k . It is clear that if $\|e_{k-1}\| < \frac{1}{2}\Delta_S$,

$$\left| P_k^{\text{req}} - P_k^{\text{imp}} \right| = \left| P_k^{\text{req}} - \text{proj}_{\mathcal{S}_{i_k}}(P_k^{\text{req}} - e_{k-1}^P) \right| < \Delta_S$$

for any projection operator. The boundary case is when $\|e_{k-1}\| = \frac{1}{2}\Delta_S$ and $P_k^{\text{req}} \in \mathcal{S}_{i_k}$. In such a case, as the ties are broken towards P_k^{req} when using the operator of Definition 9, it holds that $|P_k^{\text{req}} - P_k^{\text{imp}}| = 0$. Hence, $|P_k^{\text{req}} - P_k^{\text{imp}}| < \Delta_S$ for any k .

To conclude the proof of property (i), consider any algorithm that achieves the bounded accumulated-error property, and set $\varepsilon > 0$. Let $i^* \in [n]$ be arbitrary and let $i_k := i^*$ for every $k \in \mathbb{N}_{>0}$. Let $P_1, P_2 \in \mathcal{S}_{i^*}$ such that $P_2 = P_1 + \Delta_S$. (As $\Delta_S > 0$, such choice is guaranteed to exist.) Let $P_k^{\text{req}} := P_1 + \varepsilon$ for all $k \in \mathbb{N}_{>0}$. By the assumption on the algorithm, we have that

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K P_k^{\text{imp}} = P_1 + \varepsilon.$$

Hence, there exists k_o such that $P_{k_o}^{\text{imp}} \geq P_2$, and $P_{k_o}^{\text{imp}} - P_{k_o}^{\text{req}} \geq P_2 - P_1 - \varepsilon = \Delta_S - \varepsilon$.

Property (ii) follows then directly from property (i): If $P_{k_o}^{\text{req}} \in \mathcal{S}_{i_{k_o}}$ holds for some k_o , and $|P_{k_o}^{\text{imp}} - P_{k_o}^{\text{req}}| < \Delta_S$, then by definition of the stepsize of \mathcal{S}_{i_k} in the uniform case, we have that $P_{k_o}^{\text{imp}} = P_{k_o}^{\text{req}}$. ■

C. Example: Resource Agent for Heating a Building

In this section, we present a concrete resource-agent example: we will design a resource agent for managing the temperature in a building with several rooms. The reason for showing this example is twofold. First, we wish to give a concrete example of a resource agent that controls a load that can only implement power setpoints from a discrete set. Second, the resource-agent design shows a concrete usage example of the COMMELEC framework, and might serve as a basis for an actual resource-agent implementation.

The heating system's objective is to keep the rooms' temperatures within a certain range. For rooms whose temperature lies in that range, there is some freedom in the choice of the control actions related to those rooms. The resource agent's job is to monitor the building and spot such degrees of freedom, and expose them to the grid agent, which can then exploit those for performing *Demand Response*.

Our example is inspired by [7], which also considers the problem of controlling the temperature in the rooms of a building using multiple heaters. We address two issues that were not addressed in [7]:

- 1) We show that by rounding requested setpoints into implementable setpoints using (14) (where F_k is a suitably defined quantizer) we obtain a resource agent with bounded accumulated-error.
- 2) We prevent the heaters from switching on and off with the same frequency as COMMELEC's control frequency, which is crucial in an actual implementation.

1) Simple Case: a Single Heater: For simplicity, we first analyze a scenario with only one heater. The main aspects of our proposed design (as mentioned above) are in fact independent of the number of heaters, and we think that those aspects are more easily understood in this simple case. We will generalize our example to an arbitrary number of heaters in Section V-C2.

a) Model and Intended Behavior: We model the heater as a purely resistive load (it does not consume reactive power) that can be either active ("on") or inactive ("off"). It consumes $P_{\text{heat}} > 0$ Watts while being active, and zero Watts while being inactive.

From the perspective of the resource agent, the heater has a *state* that consists of two binary variables: $s_k \in \{0, 1\}$, which corresponds to whether the heater is on ($s_k = 1$) or off ($s_k = 0$), and $\ell_k \in \{0, 1\}$ indicates whether the heater is "locked", in which case we cannot switch on or switch off the heater. Formally, if $\ell_k = 1$ then $s_{k+1} = s_k$ necessarily holds. Hence, ℓ_k exposes a physical constraint of the heater, namely that it cannot (or should not) be switched on and off with arbitrarily high frequency. In the typical case where the minimum switching period of the heater is (much) larger than the COMMELEC's control period (≈ 100 ms), the heater will "lock" immediately after a switch, i.e., assuming $\ell_k = 0$, setting s_{k+1} such that $s_{k+1} \neq s_k$ will induce $\ell_{k'} = 1$ for every $k' \in [k+1, k+K]$, after which $\ell_{k+K+1} = 0$. Here, $K \in \mathbb{N}$ represents the minimum number of timesteps for which the heater cannot change its state from on to off or *vice versa*.

Suppose that the heater is placed in a room, and that the temperature of this room is a scalar quantity. (We do not aim here to model heat convection through the room or anything like that.) The temperature in the room, denoted as T_k , should remain within predefined "comfort" bounds,

$$T_k \in [T_{\min}, T_{\max}] \quad k = 1, 2, \dots$$

where $T_{\min}, T_{\max} \in \mathbb{R}$. If T_k is outside this interval (and only if $\ell_k = 0$), then the resource agent should take the trivial action, i.e., ensure that the heater is active if $T_k < T_{\min}$, and inactive if $T_k > T_{\max}$. The more interesting case is if T_k lies in $[T_{\min}, T_{\max}]$ (again, provided that $\ell_k = 0$), as this gives rise to some *flexibility* in the heating system: the degree of freedom here is whether to switch the heater on or off, which obviously directly corresponds to the total power consumed by the heating system. The goal is to delegate this choice to the grid agent, which we can accomplish by defining an appropriate COMMELEC advertisement.

b) Defining the Advertisement and the Rounding Behavior: Let the discrete set of implementable real-power setpoints at time k be defined as

$$\mathcal{I}_k := \begin{cases} \{0\} & \text{if } (\ell_k = 0 \wedge T_k > T_{\max}) \vee \\ & (\ell_k = 1 \wedge s_k = 0), \\ \{-P_{\text{heat}}, 0\} & \text{if } \ell_k = 0 \wedge T_{\min} \leq T_k \leq T_{\max}, \\ \{-P_{\text{heat}}\} & \text{if } (\ell_k = 0 \wedge T_k < T_{\min}) \vee \\ & (\ell_k = 1 \wedge s_k = 1), \end{cases}$$

where \wedge and \vee stand for "and" and "or", respectively. Note that \mathcal{I}_k only contains non-positive numbers, by the convention

in COMMELEC that *consuming* real power corresponds to *negative* values for P . We define the PQ profile as in Theorem 2, with \mathcal{S}_{i_k} (as appearing in Theorem 2) equal to \mathcal{I}_k , i.e.,

$$\mathcal{A}_k := \text{conv}(\mathcal{I}_k \times \{0\}) \subset \mathbb{R}^2, \quad \forall k \in \mathbb{N}_{>0}.$$

We adopt (11) as the rule to compute u_k^{imp} . As can be seen from (11), the relation between the implemented setpoint u_k^{imp} and u_k^{req} is deterministic. We can expose this relation to the grid agent by means of a belief function (valid for timestep k),

$$\begin{aligned} \text{BF}_k : \quad \mathcal{A}_k &\rightarrow \mathcal{P}(\mathbb{R}^2) \\ (p, q) &\mapsto \{(\text{proj}_{\mathcal{I}_k}(p - e_{k-1}^P), 0)\}. \end{aligned}$$

We leave the choice of a cost function to the designer of an actual resource agent, because: a) our theorems and the above resource-agent design are independent of this choice, and b) such choice typically depends on scenario-specific details.

c) *Upper Bound on the Accumulated Error:* From Theorem 2 we immediately get the following corollary.

Corollary 1. *The single-heater resource agent as defined in Section V-C1b has $\frac{1}{2}P_{\text{heat}}$ -bounded accumulated-error.*

Remark 1. *The bound given in Corollary 1 is tight in the sense that we can construct hypothetical cases in our single-heater scenario for which $|e_k| = |e_k^P| = \frac{1}{2}P_{\text{heat}}$ for some k . For example, take $k = 1$, suppose that $T_{\min} \leq T_1 \leq T_{\max}$ so that $\mathcal{I}_1 = \{-P_{\text{heat}}, 0\}$ and let $P_1^{\text{req}} = -\frac{1}{2}P_{\text{heat}}$. It then follows that $P_1^{\text{imp}} = \text{proj}_{\mathcal{I}_1}(P_1^{\text{req}} - e_0^P) = \text{proj}_{\mathcal{I}_1}(P_1^{\text{req}}) \in \mathcal{I}_1$, which gives $|e_1^P| = |P_1^{\text{imp}} - P_1^{\text{req}}| = \frac{1}{2}P_{\text{heat}}$.*

2) *General Case: an Arbitrary Number of Heaters:* Here, we extend the single-heater case to a setting with r heaters, for $r \in \mathbb{N}, r \geq 1$ arbitrary. As we will see, also this multi-heater case can be analyzed using the tools introduced at the beginning of Section V.

Like in the single-heater case, we assume that each heater is purely resistive. We furthermore assume that heater i consumes P_i^{heat} Watts of power when active (and zero power when inactive), for every $i \in [r]$. Also similarly to the single-heater case, we assume that each heater is placed in a separate room, whose (scalar) temperature is denoted as $T_k^{(i)}$. Not surprisingly, our objective shall now be to keep the temperature in each room within the predefined comfort bounds, i.e.,

$$T_k^{(i)} \in [T_{\min}, T_{\max}] \quad \forall k \in \mathbb{N}_{>0}, \forall i \in [r].$$

In the one-heater case, the only degree of freedom that can be present is the choice to switch that heater on or off. In case of multiple heaters, there is potentially some freedom in choosing which subset of the heaters to activate, and note that there will typically² be an exponential number of those subsets (exponential in the number of heaters). Each subset corresponds to a certain total power consumption, i.e., a power setpoint. As in the single-heater case, the PQ profile will be defined as the convex hull of the collection of these setpoints. When the grid agent requests some setpoint from the PQ profile, the resource agent has to select an

appropriate subset whose corresponding setpoint is closest (in the Euclidean sense) to the requested setpoint. Note that there can be several subsets of heaters that correspond to the same setpoint. A simple method to resolve this ambiguity would be, for example, to choose the subset consisting of the coldest rooms, however, as this topic is beyond the scope of this work, we leave the choice of such a selection method to the resource-agent designer.

a) *Characterizing the Set of Implementable Setpoints:*

When going from the single-heater setting to a multiple-heaters scenario, we merely need to re-define \mathcal{I}_k , which we will name $\tilde{\mathcal{I}}_k$ here to avoid confusion with the single-heater case. The definition of the PQ profile, belief function, and rule for computing u_k^{imp} given in Section V-C1b also apply to the multi-heater case, provided that all occurrences of \mathcal{I}_k in those definitions are replaced by $\tilde{\mathcal{I}}_k$.

For every $i \in [r]$, let $s_k^{(i)}$ and $\ell_k^{(i)}$ represent the state variables s_k and ℓ_k (as defined in the single-heater case) for the i -th heater. Let $\mathcal{L}_k := \{i \in [r] : \ell_k^{(i)} = 1\}$ denote the set of rooms whose heater is locked at timestep k . Furthermore, let $\mathcal{C}_k := \{i \in [r] : T_k^{(i)} < T_{\min}\}$ and $\mathcal{W}_k := \{i \in [r] : T_{\min} \leq T_k^{(i)} \leq T_{\max}\}$. Informally speaking, \mathcal{C}_k contains the rooms that are “too cold”, and \mathcal{W}_k the rooms whose temperatures are within the comfort bounds.

If $A \subseteq [r]$, we write \bar{A} for the complement with respect to $[r]$, i.e. $\bar{A} := [r] \setminus A$.

Let

$$\tilde{\mathcal{I}}_k := \left\{ a_k - \sum_{i \in \mathcal{S}} P_i^{\text{heat}} : \mathcal{S} \subseteq \bar{\mathcal{L}}_k \cap \mathcal{W}_k \right\} \quad (12)$$

represent the set of implementable (active) power setpoints, with

$$a_k := - \sum_{i \in \mathcal{L}_k} s_k^{(i)} P_i^{\text{heat}} - \sum_{j \in \mathcal{L}_k \cap \mathcal{C}_k} P_j^{\text{heat}}.$$

b) *Upper Bound on the Accumulated Error:* As in the one-heater example, we use Theorem 2 to bound the accumulated error of the resource agent. To this end, let $\tilde{\mathcal{S}}$ denote the collection of all possible sets $\tilde{\mathcal{I}}_k$ (12). It is easy to see that this is a finite collection. Further, the maximum stepsize of this collection (Definition 7) is given by $\Delta_{\tilde{\mathcal{S}}} = \max_{i \in [r]} P_i^{\text{heat}}$. This gives us the following corollary.

Corollary 2. *The multiple-heaters resource agent has $(\frac{1}{2} \max_{i \in [r]} P_i^{\text{heat}})$ -bounded accumulated-error.*

VI. UNCERTAIN RESOURCES

In this section, we show how to achieve boundedness of the accumulated error for *uncertain* resources, as given in Definition 4. In particular, we focus on resources that are affected by Nature, and hence the relation between the advertised PQ profile and the set of implementable setpoints is uncertain. This covers such resources as PV panels, wind farms, and partially controllable loads.

A. Key Property: Projection-Translation Invariance

In order to state our result, we first introduce an important definition that we use in the following.

²Provided that not too many heaters are locked.

Definition 10 (Projection-Translation Invariance). Fix $d \in \mathbb{N}$. Let $\mathcal{D} \subseteq \mathbb{R}^d$ be a given convex compact set. We say that a convex set $\mathcal{I} \subseteq \mathbb{R}^d$ is a projection-translation invariant subset of \mathcal{D} if $\mathcal{I} \subseteq \mathcal{D}$ and for every $v \in \mathcal{D}$ and $u \in \mathcal{I}$, it holds that

$$u + v - \text{proj}_{\mathcal{I}}(v) \in \mathcal{D}.$$

A helpful interpretation of Definition 10 is to view $\tau := v - \text{proj}_{\mathcal{I}}(v)$ as a translation vector in the projection direction of v to \mathcal{I} . Then, projection-translation invariance guarantees that the translation of \mathcal{I} over τ remains contained in \mathcal{D} .

It is easy to see that in the one dimensional case ($d = 1$), projection-translation invariance is satisfied for intervals.

Proposition 2. For any $a \leq c \leq d \leq b$, the interval $\mathcal{I} = [c, d]$ is a projection-translation invariant subset of the interval $\mathcal{D} = [a, b]$.

Proof: Let $v \in \mathcal{D}$ and $u \in \mathcal{I}$. If $v \in \mathcal{I}$, trivially $u + v - \text{proj}_{\mathcal{I}}(v) = u \in \mathcal{D}$. Now, if $v > d$, we have that

$$u + v - \text{proj}_{\mathcal{I}}(v) \leq d + b - d = b, \text{ and}$$

$$u + v - \text{proj}_{\mathcal{I}}(v) \geq c + a - d \geq a,$$

namely $u + v - \text{proj}_{\mathcal{I}}(v) \in \mathcal{D}$. Similarly, if $v < c$, it holds that

$$u + v - \text{proj}_{\mathcal{I}}(v) \leq d + b - c \leq b, \text{ and}$$

$$u + v - \text{proj}_{\mathcal{I}}(v) \geq c + a - c = a,$$

namely $u + v - \text{proj}_{\mathcal{I}}(v) \in \mathcal{D}$. \blacksquare

For $d = 2$ and if \mathcal{I} is a convex polygon, then we can construct a collection of sets $\mathbb{D}_{\mathcal{I}}$ such that \mathcal{I} is guaranteed (by construction) to be a projection-translation-invariant subset of every $\mathcal{D} \in \mathbb{D}_{\mathcal{I}}$ (see Construction 1 in Section VII-D for details).

Constructing sets with the projection-translation invariant property in higher dimensions is left for further work. It might well be that our construction for polygons carries over to polytopes in \mathbb{R}^d for arbitrary d .

B. Main Results

The following result provides an algorithm for uncertain resources that are characterized by the projection-translation invariance property. The proof is deferred to Section VII.

Theorem 4. Consider an uncertain resource as per Definition 4. In particular, the set of implementable (feasible) setpoints at time step $k \in \mathbb{N}_{>0}$ is given by a convex set \mathcal{I}_k that is not known at the time of advertising \mathcal{A}_k . Suppose in addition that, for every k , \mathcal{I}_k is a projection-translation invariant subset of a given convex compact set $\mathcal{D} \subseteq \mathbb{R}^2$. For every $k \in \mathbb{N}$, let $e_k \in \mathbb{R}$ be the accumulated error as defined in Eqn. (2). Then, if the resource agent

- uses a persistent predictor to advertise the PQ profile, namely sends $\mathcal{A}_k = \mathcal{I}_{k-1}$; and
- implements $u_k^{\text{imp}} = \text{proj}_{\mathcal{I}_k}(u_k^{\text{req}} - e_{k-1})$ for every $k \in \mathbb{N}_{>0}$,

then it has $(\text{diam } \mathcal{D})$ -bounded accumulated-error, where $\text{diam } \mathcal{D}$ is the diameter of \mathcal{D} .

Note that Theorem 4 uses the projection-translation-invariance property in a “many-to-one” relation: many sets \mathcal{I}_k are required to be projection-translation-invariant subsets of one set \mathcal{D} . The theorem does not say how the set \mathcal{D} can be constructed; it simply supposes that such a set \mathcal{D} is “given”.

In Section VI-C, we will see an example where the sets \mathcal{I}_k are selected from a particular parameterized collection of sets for which it turns out to be straightforward to find a set \mathcal{D} with the required property. And, in Section VII-D, we deal with the problem of finding such \mathcal{D} for a more general case.

In the one-dimensional case, it is trivial to construct a set \mathcal{D} with the required “many-to-one” property: using Proposition 2, it is easy to see that such a \mathcal{D} can be constructed by taking the union of all intervals \mathcal{I}_k ; we will use this in the following corollary.

Corollary 3. Consider an uncertain real-power only resource. In particular, the set of implementable (feasible) setpoints at time step $k \in \mathbb{N}_{>0}$ is given by an interval \mathcal{I}_k that is not known at the time of advertising \mathcal{A}_k . Then, if the resource agent

- uses a persistent predictor to advertise the PQ profile, namely sends $\mathcal{A}_k^P = \mathcal{I}_{k-1}$ ($\mathcal{A}_k^Q = \{0\}$); and
- implements $P_k^{\text{imp}} = \text{proj}_{\mathcal{I}_k}(P_k^{\text{req}} - e_{k-1}^P)$ for every $k \in \mathbb{N}_{>0}$,

then it has $(\text{diam } \mathcal{D})$ -bounded accumulated-error, where $\mathcal{D} := \text{conv}(\bigcup_{k=1}^{\infty} \mathcal{I}_k)$.

C. Example: Photovoltaic (PV) System

Here, we explain how we can apply Theorem 4 to devise a resource agent for a PV system with bounded accumulated-error.

Let S_{rated} and ϕ_{max} denote the rated power of the converter and angle corresponding to the minimum power factor, respectively. We suppose that these quantities are given (they correspond to physical properties of the PV system), and that $S_{\text{rated}} \geq 0$ and $0 \leq \phi_{\text{max}} < \pi$. Note that for any power setpoint (P, Q) , the rated power imposes the constraint $P^2 + Q^2 \leq S_{\text{rated}}^2$; the angle ϕ_{max} imposes that $\arctan(Q/P) \leq \phi_{\text{max}}$.

Let us now choose $P_{\text{max}}, \varphi \in \mathbb{R}$ such that $0 \leq P_{\text{max}} \leq S_{\text{rated}}$, $0 \leq \varphi \leq \phi_{\text{max}}$, and $\frac{P_{\text{max}}}{\cos \varphi} = S_{\text{rated}}$, and let

$$\mathcal{T}(x) := \{(P, Q) \in \mathbb{R}^2 : 0 \leq P \leq x, \frac{|Q|}{P} \leq \tan \varphi\}$$

be a triangle-shaped set in the PQ plane; see Figure 3 for an illustration. Note that for any combination of P_{max} and φ , the triangle $\mathcal{T}(x)$ for any $x \in [0, P_{\text{max}}]$ is fully contained in the disk that corresponds to the rated-power constraint, and, moreover, the two upper corner points of $\mathcal{T}(P_{\text{max}})$ lie on the boundary of that disk.

Let p_k^{max} be the maximum real power available at timestep $k \in \mathbb{N}$ (typically determined by the solar irradiance). Using p_k^{max} , we define the set of implementable points at timestep k as

$$\mathcal{I}_k := \mathcal{T}(\min(p_k^{\text{max}}, P_{\text{max}})).$$




Fig. 3. The parameterized collection of sets $\{\mathcal{T}(x) : x \in [0, P_{\max}]\}$, and its relation to the rated-power constraint of the PV converter.

Corollary 4. Let P_{\max} , φ , and \mathcal{I}_k (for every $k \in \mathbb{N}$) be defined as above for a given PV system. A resource agent for this PV system that advertises $\mathcal{A}_k := \mathcal{I}_{k-1}$ and implements $u_k^{\text{imp}} = \text{proj}_{\mathcal{I}_k}(u_k^{\text{req}} - e_{k-1})$ for every $k \in \mathbb{N}_{>0}$ has c -bounded accumulated-error with

$$c := \max \left\{ \frac{P_{\max}}{\cos \varphi}, 2 P_{\max} \tan \varphi \right\}.$$

The main ingredient of the proof of Corollary 4 is the following lemma, whose proof is given in Section VII-E.

Lemma 3. Every set in the collection $\{\mathcal{T}(x) : x \in [0, P_{\max}]\}$ is projection-translation-invariant with respect to $\mathcal{D} := \mathcal{T}(P_{\max})$.

Proof of Corollary 4.: By definition of \mathcal{I}_k and by Lemma 3, each member of the sequence $(\mathcal{I}_k)_k$ is projection-translation-invariant with respect to $\mathcal{T}(P_{\max})$. Because the resource agent computes the PQ profile \mathcal{A}_k and u_k^{imp} in accordance with the requirements of Theorem 4, we can apply the latter to conclude that the resource agent has $\text{diam} \mathcal{T}(P_{\max})$ -bounded accumulated-error. Since $\mathcal{T}(P_{\max})$ is an isosceles triangle, its diameter is either $P_{\max}/\cos \varphi$ (the length of one of its legs) for $\varphi \leq \frac{\pi}{6}$ or $2 P_{\max} \tan \varphi$ otherwise (the length of its base). Hence, the claim follows. ■

VII. ACHIEVING BOUNDED ACCUMULATED-ERROR THROUGH ERROR DIFFUSION: GENERIC APPROACH AND PROOFS

In this section, we propose a generic approach for achieving bounded accumulated-error via the error-diffusion method. The results in this section are used to prove the results of Sections V and VI; however, they can also be used to design agents for other resources that satisfy appropriate conditions. The reader that is interested in the performance of the RAs proposed in Sections V and VI can safely skip the present section and proceed to Section VIII.

We next introduce two general concepts: the first (Definition 11) is a generalization of the c -bounded accumulated-error concept, and the second (Definition 12) is a generalization of the projection concept.

Definition 11. For any set $\mathcal{G} \in \mathcal{B}(\mathbb{R}^2)$, we say that a given resource agent has \mathcal{G} -contained accumulated-error if

$$e_k \in \mathcal{G}, \quad k \geq 1.$$

Note that having \mathcal{G} -contained accumulated-error implies having c -bounded accumulated-error (Definition 5) with $c \geq \max\{\|x\| : x \in \mathcal{G}\}$.

Definition 12. Fix $d \in \mathbb{N}$. Let $F : \mathcal{D} \rightarrow \mathcal{I}$ be an arbitrary map, where $\mathcal{D}, \mathcal{I} \subseteq \mathbb{R}^d$. We say that F is a \mathcal{G} -approximation if there exists a set $\mathcal{G} \in \mathcal{B}(\mathbb{R}^d)$ such that

$$F(x) - x \in \mathcal{G}$$

holds for all $x \in \mathcal{D}$.

Note that the projection operator is a \mathcal{G} -approximation for suitably chosen \mathcal{G} .

The following lemma is a key result that explains why the error-diffusion method leads to bounded accumulated-error.

Lemma 4. Let $\mathcal{G} \in \mathcal{B}(\mathbb{R}^2)$ and let $k \in \mathbb{N}_{>0}$ be arbitrary. Let $\mathcal{D}_k, \mathcal{I}_k \subseteq \mathbb{R}^2$ be non-empty sets and let $F_k : \mathcal{D}_k \rightarrow \mathcal{I}_k$ be a \mathcal{G} -approximation. Let $e_k \in \mathbb{R}^2$ be the accumulated error as defined in Eqn. (2). Let $u_k^{\text{req}} \in \mathbb{R}^2$. Then, if

$$u_k^{\text{req}} - e_{k-1} \in \mathcal{D}_k \quad (13)$$

holds, and

$$u_k^{\text{imp}} := F_k(u_k^{\text{req}} - e_{k-1}), \quad (14)$$

we have that

$$e_k \in \mathcal{G}.$$

Proof: Note that we can write e_k in the following equivalent recursive form:

$$\begin{aligned} e_k &= e_{k-1} + u_k^{\text{imp}} - u_k^{\text{req}}, \\ &= e_{k-1} + F_k(u_k^{\text{req}} - e_{k-1}) - u_k^{\text{req}} \quad (\text{using Eqn. (14)}) \\ &= F_k(u_k^{\text{req}} - e_{k-1}) - (u_k^{\text{req}} - e_{k-1}). \end{aligned} \quad (15)$$

The result follows since F_k is a \mathcal{G} -approximation. ■

Now, the main idea of our approach is as follows: if (i) a resource agent can be characterized by a sequence of \mathcal{G} -approximations $F_k, k \in \mathbb{N}_{>0}$ for some $\mathcal{G} \in \mathcal{B}(\mathbb{R}^2)$ that is independent of k , where this sequence of maps $\{F_k\}$ represents the physics of the controlled resource as well as any internal (low-level) control loops in the resource itself or in the resource agent, and (ii) u_k^{imp} is given by Eqn. (14) for every $k \in \mathbb{N}_{>0}$, then that resource agent achieves \mathcal{G} -contained accumulated-error if we can invoke Lemma 4 for every $k \in \mathbb{N}_{>0}$. Conditioned on (i) and (ii), we may invoke Lemma 4 for every k if we can prove that condition (13) is satisfied for all k . In the following subsections, we show how we can ensure (13) for deterministic and uncertain resources.

A. Result for Deterministic Resources

In this section, we focus on deterministic resources as per Definition 4. In practice this means that the set of implementable setpoints at time step k is known in advance, hence can be used to advertise the PQ profile \mathcal{A}_k .

Theorem 5. Let $\mathcal{G} \in \mathcal{B}(\mathbb{R}^2)$. Let $e_k \in \mathbb{R}^2$ be the accumulated error as defined in Eqn. (2) for every $k \in \mathbb{N}$. For every $k \in \mathbb{N}_{>0}$, let $\mathcal{I}_k \in \mathcal{B}(\mathbb{R}^2)$, $\mathcal{A}_k := \text{conv}(\mathcal{I}_k)$ and $u_k^{\text{req}} \in \mathcal{A}_k$, let

$\mathcal{D}_k := \mathcal{A}_k + \mathcal{G}$, and let $F_k : \mathcal{D}_k \rightarrow \mathcal{I}_k$ a \mathcal{G} -approximation. Then, if a resource agent implements $u_k^{\text{imp}} = F_k(u_k^{\text{req}} - e_{k-1})$ for every $k \in \mathbb{N}_{>0}$, it has \mathcal{G} -contained accumulated-error.

The trick here is to define \mathcal{D}_k (the domain of F_k) as the “ \mathcal{G} -inflation” of \mathcal{A}_k , such that condition (13) is trivially satisfied. Of course, to use this theorem for an actual application, it remains to ensure that the (application-specific) map F_k is a \mathcal{G} -approximation on this particular domain. In Section VII-B, we prove that for the discrete resources of Section V, the corresponding map has this required property.

Proof: We will prove the statement using induction on k . By definition of e_k , the statement holds for $k = 0$. Suppose that the statement holds for timestep $k - 1$ (for arbitrary $k \in \mathbb{N}_{>0}$). Since $u_k^{\text{req}} \in \mathcal{A}_k$ and because $e_{k-1} \in \mathcal{G}$ (by the induction hypothesis), it holds that

$$u_k^{\text{req}} - e_{k-1} \in \mathcal{A}_k + \mathcal{G} = \mathcal{D}_k$$

by construction. Furthermore, F_k is a \mathcal{G} -approximation and $u_k^{\text{imp}} = F_k(u_k^{\text{req}} - e_{k-1})$, hence we may invoke Lemma 4 to conclude that $e_k \in \mathcal{G}$. ■

Remark 2. From the proof, it might seem that we do not require any relation between \mathcal{A}_k and \mathcal{I}_k , however, without any such relation it will be impossible to construct F_k such that it is a \mathcal{G} -approximation for every k . We define $\mathcal{A}_k := \text{conv}(\mathcal{I}_k)$ as it satisfies our needs, but in principle one could be more general here, and define \mathcal{A}_k as an arbitrary function of \mathcal{I}_k .

B. Proof of Theorem 2

Fact 1. The bound

$$|\text{proj}_{\mathcal{S}}(x) - x| \leq \frac{\Delta_{\mathcal{S}}}{2}$$

holds for every finite non-empty set $\mathcal{S} \subset \mathbb{R}$ and for every $x \in [\min \mathcal{S}, \max \mathcal{S}]$.

Lemma 5. Let \mathbb{S} be a finite non-empty collection of sets $\mathbb{S} := \{\mathcal{S}_i\}_{i \in [n]}$ where $n \in \mathbb{N}_{>0}$ and $\mathcal{S}_i \subset \mathbb{R}$ is a finite non-empty set for all $i \in [n]$. Let $\Delta_{\mathbb{S}}$ denote the maximum stepsize of \mathbb{S} , and define $\mathcal{G} := [-\frac{1}{2}\Delta_{\mathbb{S}}, \frac{1}{2}\Delta_{\mathbb{S}}] \times \{0\}$. Let $(i_k)_{k \in \mathbb{N}_{>0}}$ be any sequence with $i_k \in [n]$, and let $\mathcal{I}_k := \mathcal{S}_{i_k} \times \{0\}$, $\mathcal{D}_k := \text{conv}(\mathcal{I}_k) + \mathcal{G}$ and

$$\begin{aligned} F_k : \mathcal{D}_k &\rightarrow \mathcal{I}_k \\ (p, q) &\mapsto (\text{proj}_{\mathcal{S}_{i_k}}(p), 0) \end{aligned}$$

be defined for every $k \in \mathbb{N}_{>0}$. Then F_k is a \mathcal{G} -approximation for all $k \in \mathbb{N}_{>0}$.

Proof: We need to show that $F_k(x) - x \in \mathcal{G}$ for every $x \in \mathcal{D}_k$ and every $k \in \mathbb{N}_{>0}$. Since $\mathcal{D}_k^Q = \mathcal{G}^Q = 0$, it suffices to show that $\text{proj}_{\mathcal{I}_k}(y) - y \in \mathcal{G}^P = [-\frac{1}{2}\Delta_{\mathbb{S}}, \frac{1}{2}\Delta_{\mathbb{S}}]$ for every $y \in \mathcal{D}_k^P = [\min \mathcal{I}_k - \frac{1}{2}\Delta_{\mathbb{S}}, \max \mathcal{I}_k + \frac{1}{2}\Delta_{\mathbb{S}}]$ and every $k \in \mathbb{N}_{>0}$.

Let k be arbitrary. If $y \in [\min \mathcal{I}_k, \max \mathcal{I}_k]$, then

$$|\text{proj}_{\mathcal{I}_k}(y) - y| \leq \frac{\Delta_{\mathcal{S}_{i_k}}}{2} \leq \frac{\Delta_{\mathbb{S}}}{2}$$

by Fact 1 and by Definition 7. Otherwise, i.e., if $y \notin \text{conv}(\mathcal{I}_k)$, y will be projected to the closest boundary of $\text{conv}(\mathcal{I}_k)$, for which $|\text{proj}_{\mathcal{I}_k}(y) - y| \leq \frac{1}{2}\Delta_{\mathbb{S}}$ holds by definition of \mathcal{D}_k . ■

Proof of Theorem 2: Lemma 5 guarantees that the map

$$F_k(u_k^{\text{req}}) := (\text{proj}_{\mathcal{S}_{i_k}}(P_k^{\text{req}} - e_{k-1}^P), 0)$$

is a \mathcal{G} -approximation on the domain $\mathcal{A}_k + \mathcal{G}$ where $\mathcal{A}_k := \text{conv}(\mathcal{I}_k)$, $\mathcal{I}_k := \text{conv}(\mathcal{S}_{i_k} \times \{0\})$ and $\mathcal{G} := [-\frac{1}{2}\Delta_{\mathbb{S}}, \frac{1}{2}\Delta_{\mathbb{S}}] \times \{0\}$. Because the implemented setpoint is computed as

$$u_k^{\text{imp}} = (\text{proj}_{\mathcal{S}_{i_k}}(P_k^{\text{req}} - e_{k-1}^P), 0)$$

for every $k \in \mathbb{N}_{>0}$, we can apply Theorem 5 to conclude that the resource agent has \mathcal{G} -contained accumulated-error, which implies that the resource agent has $\frac{1}{2}\Delta_{\mathbb{S}}$ -bounded accumulated-error by definition of \mathcal{G} . ■

C. Proof of Theorem 4

Proof: Consider the mapping $F_k : \mathcal{D} \rightarrow \mathcal{I}_k$, $x \mapsto \text{proj}_{\mathcal{I}_k}(x)$. Recall that \mathcal{I}_k is, in particular, a convex subset of \mathcal{D} by the definition of projection-translation invariance. Thus, $\|F_k(x) - x\| \leq \text{diam } \mathcal{D}$ holds for all $x \in \mathcal{D}$ and all $k \in \mathbb{N}_{>0}$; namely, F_k is a $(\text{diam } \mathcal{D})$ -approximation for all $k \in \mathbb{N}_{>0}$.

Next, we prove that (i) $u_k^{\text{req}} - e_{k-1} \in \mathcal{D}$ and that (ii) $\|e_k\| \leq \text{diam } \mathcal{D}$ holds for all $k \in \mathbb{N}_{>0}$, using induction. Clearly, claim (i) holds for $k = 1$ as $u_1^{\text{req}} - e_0 = u_1^{\text{req}} \in \mathcal{A}_1 = \mathcal{I}_0 \subseteq \mathcal{D}$. This result allows us to invoke Lemma 4 and conclude that (ii) also holds for $k = 1$.

Now suppose that $u_k^{\text{req}} - e_{k-1} \in \mathcal{D}$ holds for arbitrary $k \in \mathbb{N}_{>0}$ (the induction hypothesis). For $k + 1$, we have, as in (15),

$$u_{k+1}^{\text{req}} - e_k = u_{k+1}^{\text{req}} + (u_k^{\text{req}} - e_{k-1}) - \text{proj}_{\mathcal{I}_k}(u_k^{\text{req}} - e_{k-1}).$$

Observe that $u_{k+1}^{\text{req}} \in \mathcal{I}_k$ as the resource agent uses the persistent predictor for the PQ profile. Also, by the induction hypothesis, $u_k^{\text{req}} - e_{k-1} \in \mathcal{D}$. Hence, invoking the projection-translation invariance property of \mathcal{I}_k (Definition 10) with $u \equiv u_{k+1}^{\text{req}}$ and $v \equiv u_k^{\text{req}} - e_{k-1}$, we obtain that $u_{k+1}^{\text{req}} - e_k \in \mathcal{D}$, which proves (i) for $k + 1$. Again, this allows us to invoke Lemma 4 for $k + 1$ and conclude that $\|e_{k+1}\| \leq \text{diam } \mathcal{D}$ which completes the induction argument, and with that, the proof. ■

D. Projection-Translation Invariance: Constructing Supersets of Polygons

In this section we present, for a given convex polygon \mathcal{I} , how to construct a superset \mathcal{D} such that \mathcal{I} is a projection-translation invariant subset of \mathcal{D} . The superset is not unique; in fact, the degrees of freedom in constructing the superset give rise to an (infinite) collection of supersets, $\mathbb{D}_{\mathcal{I}}$, which we have already introduced in Section VI-A.

We will use the construction in the following way: when given a collection of sets $\{\mathcal{I}_k\}_k$, and if, for every k , \mathcal{I}_k is a convex polygon, then we can construct a *minimal* set \mathcal{D} with the “many-to-one” property (which means that for every k , \mathcal{I}_k is a projection-translation invariant subset of \mathcal{D}) as

$$\mathcal{D} \in \arg \min \{ \text{diam } \mathcal{W} : \mathcal{W} \in \bigcap_k \mathbb{D}_{\mathcal{I}_k} \}, \quad (16)$$

where we claim that the intersection in (16) is always non-empty and always contains a set with bounded diameter, provided that all sets \mathcal{I}_k are bounded.




Fig. 4. The cones $C := \text{cone}(v, w)$, $-C$ and C^o , and the half-space $\mathcal{H}(a, x_0)$.

For our construction, we need the following definitions. (See also Figure 4 for some pictorial examples of these definitions.)

Definition 13 (Cone). We define the (convex) cone of any two vectors $v, w \in \mathbb{R}^2$ as

$$\text{cone}(v, w) := \{\alpha v + \beta w : \alpha, \beta \in \mathbb{R}, \alpha \geq 0, \beta \geq 0\}.$$

Definition 14 (Polar Cone). For any cone $C \subset \mathbb{R}^2$, we define the polar cone of C as

$$C^o := \{y \in \mathbb{R}^2 : y^T x \leq 0 \quad \forall x \in C\}$$

Definition 15 (Half Space). Let

$$\begin{aligned} \mathcal{H} : \mathbb{R}^2 \times \mathbb{R}^2 &\rightarrow \mathcal{P}(\mathbb{R}^2) \\ (a, x_0) &\mapsto \{x \in \mathbb{R}^2 \mid a^T(x - x_0) \leq 0\}. \end{aligned}$$

Note that $\mathcal{H}(a, x_0)$ is a closed halfspace that extends in direction $-a$ and where x_0 lies on its associated halfplane.

Our construction is given below. See Figure 5 for a graphical example.

Construction 1. Let $\mathcal{I} \subset \mathbb{R}^2$ be a simple convex polygon. Let n denote the number of vertices of \mathcal{I} , let v_1 be an arbitrary vertex of \mathcal{I} and let v_2, \dots, v_n respectively denote the remaining vertices encountered when traversing the polygon clockwise from v_1 . Let

$$\text{wrap}(i) := \begin{cases} 1 & \text{if } i = n + 1 \\ i & \text{if } i \in [n] \\ n & \text{if } i = 0 \end{cases}$$

be defined for any $i \in \{0, \dots, n + 1\}$. Let $\mathcal{C}_i := \text{cone}(v_{\text{wrap}(i+1)} - v_i, v_{\text{wrap}(i-1)} - v_i)$ for every $i \in [n]$.

Choose an arbitrary point $\pi_1 \in (-\mathcal{C}_1) \cap \mathcal{C}_1^o$ and define $p_1 := v_1 + \pi_1$.³ Let ℓ_1 be the line through p_1 and parallel to $v_2 - v_1$. For every $j \in \{2, \dots, n - 1\}$, choose a point $\pi_j \in (\ell_{j-1} - \{v_j\}) \cap (-\mathcal{C}_j) \cap \mathcal{C}_j^o$, let $p_j := v_j + \pi_j$ and let ℓ_j be the line through p_j and parallel to $v_{j+1} - v_j$. Finally, let ℓ_n be the line through p_1 and parallel to $v_n - v_1$, and define p_n as the intersection between ℓ_{n-1} and ℓ_n .

Let $\mathcal{P}_i^{\min} := \mathcal{C}_i^o \cap \mathcal{H}(v_i - \text{proj}_{x_i}(v_i), \text{proj}_{x_i}(v_i) - v_i)$, where x_i is the line through $\text{proj}_{\ell_i}(v_i)$ and $\text{proj}_{\ell_{\text{wrap}(i-1)}}(v_i)$. Let \mathcal{R}

³If the interior angle of vertex i is *acute*, then $(-\mathcal{C}_i) \cap \mathcal{C}_i^o = -\mathcal{C}_i$. Otherwise, i.e., if the interior angle is *obtuse*, then $(-\mathcal{C}_i) \cap \mathcal{C}_i^o = \mathcal{C}_i^o$.




Fig. 5. Illustration of parts of Construction 1 for constructing a set \mathcal{D} such that \mathcal{I} is a projection-translation-invariant subset of \mathcal{D} .

be the polygon defined by $\{p_i\}_{i \in [n]}$, and let $\mathcal{P}_i^{\max} := \mathcal{C}_i^o \cap (\mathcal{R} - \{v_i\})$ for every $i \in [n]$.

Let \mathcal{G} be an arbitrary convex set such that $\bigcup_{i \in [n]} \mathcal{P}_i^{\min} \subseteq \mathcal{G} \subseteq \bigcup_{i \in [n]} \mathcal{P}_i^{\max}$. Define $\mathcal{D} := \mathcal{I} + \mathcal{G}$.

Lemma 6. For a given simple convex polygon \mathcal{I} , any set \mathcal{D} constructed using Construction 1 is such that \mathcal{I} is a projection-translation-invariant subset of \mathcal{D} .

Proof: Let v be as defined in Definition 10. First, note that the projection-translation-invariance property trivially holds when $v \in \mathcal{I}$. Hence, suppose w.l.o.g. that $v \in \mathcal{D} \setminus \mathcal{I}$. Now, observe that, by construction, the set \mathcal{G} contains all possible translation vectors $\{v - \text{proj}_{\mathcal{I}}(v)\}_{v \in \mathcal{D} \setminus \mathcal{I}}$. Namely, (i) for every $i \in [k]$ the line ℓ_i is parallel to the corresponding facet of \mathcal{I} ; (ii) $\mathcal{G} \supseteq \bigcup_{i \in [k]} \mathcal{P}_i^{\min}$ and, (iii) proj is the orthogonal projection. Hence, \mathcal{I} is a projection-translation-invariant subset of $\mathcal{D} = \mathcal{I} + \mathcal{G}$. ■

E. Proof of Lemma 3

Throughout the proof (of Lemma 3), we re-use the notation and objects that we introduced/defined in Construction 1.




Fig. 6. Invoking Construction 1 to show that $\{\mathcal{T}(x) : x \in [0, P_{\max}]\}$ is a projection-translation-invariant collection with respect to $\mathcal{D} := \mathcal{T}(P_{\max})$.

Figure 6 shows a helpful illustration of the objects that play a role in the proof.

Proof of Lemma 3.: Pick $x \in [0, P_{\max}]$ arbitrarily. We now use Construction 1 to construct \mathcal{D} from $\mathcal{T}(x)$ (where the latter corresponds to the set \mathcal{I} in Construction 1). Let $v_1 := (0, 0)$ and $n = 3$. Choose $p_1 := v_1$. Note that this choice also defines the line ℓ_1 . Choose p_2 on ℓ_1 , such that $p_2^P = P_{\max}$. This choice also defines ℓ_2 , p_3 and ℓ_3 . Let $\mathcal{G} := \bigcup_{i \in [3]} \mathcal{P}_i^{\max} = \mathcal{P}_2^{\max} \cup \mathcal{P}_3^{\max}$, since $\mathcal{P}_1^{\max} = \{(0, 0)\} \subseteq \mathcal{P}_2^{\max} \cup \mathcal{P}_3^{\max}$. Now, because $\mathcal{T}(x)$ is a convex polygon and because we followed Construction 1 to construct $\mathcal{D} = \mathcal{T}(x) + \mathcal{G}$, we can use Lemma 6 to conclude that $\mathcal{T}(x)$ is a projection-translation-invariant subset of \mathcal{D} . Because this holds true for arbitrary $x \in [0, P_{\max}]$, we may conclude that every set in the collection $\{\mathcal{T}(x) : x \in [0, P_{\max}]\}$ is projection-translation-invariant with respect to \mathcal{D} . ■

VIII. NUMERICAL ILLUSTRATION

Our aim here is to explore through numerical simulation the behavior of the closed-loop system (as introduced in Section IV) in a realistic scenario that is not covered by Hypothesis 1. Our simulation results indicate that the use of resource agents that have the bounded accumulated-error property significantly improves the behavior the closed-loop COMMELEC system.

As in [2], we take a case study that makes reference to the low voltage microgrid benchmark defined by the CIGRÉ Task Force C6.04.02 [8]. For the full description of the case study and the corresponding agents design, the reader is referred to [2].

There are two modifications compared to the original case study: (i) The PV agents are updated with the algorithm described in Section VI, and (ii) the uncontrollable load (UL2) is replaced by a controllable load (modeling a resistive heater), and the corresponding agent is implemented according to the methods described in Section V-C.

We simulate a rather extreme scenario involving a highly variable solar irradiance profile. I.e., we let the irradiance vary according to a square wave with a period of 300 ms. This will cause the PV agent's PQ profile to be highly variable, which, in turn, means that part (ii) of Hypothesis 1 will get violated frequently. Furthermore, our setup includes various resource agents that send time-variable cost functions, which violates condition (iii) of Hypothesis 1. Note that the varying state

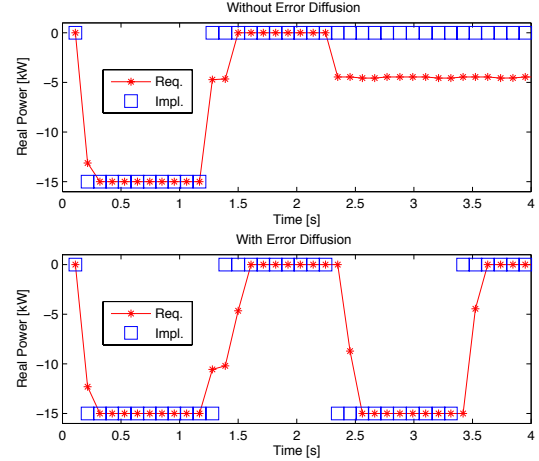


Fig. 7. Sequence of requested vs. implemented setpoints belonging to the heater agent. In the top figure (no error diffusion) the grid agent “gets stuck” on a setpoint that is not close to the optimal value. In the bottom figure (with error diffusion), the heater switches on and off with an appropriate duty cycle (the switching frequency is limited by the locking-duration parameter).

of the grid will, through the related cost term, also violate condition (iii).

We let the cost function of the PV agent be the same as in [2]; this cost function encourages to maximize active-power output. The cost function of the heater is set to a quadratic function, whose minimum lies at half the heater power, namely at -7.5 kW. With respect to the locking behavior of the heater, we let it lock for one second after a switch.

The results are shown in Figures 7–12. For comparison, we run the same scenario with resource agents for which the accumulated error might grow unboundedly. I.e., those RAs do not apply the error-diffusion technique described in this paper, instead, they just project the request to the closest implementable setpoint, like in [2]. From the results, we see the following benefits:

- Improved utilization of renewables, less curtailment. (Figure 10)
- Convergence to the optimum of the cost functions. (Figure 9)
- Delivery of the requested power on average, i.e., energy, which can be valuable for an application like a virtual power plant. (Figure 8, 9, 11 and 12)
- Less sensitivity to the choice of the grid agent's gradient-descent step size. (Figure 7)

IX. CONCLUSION

We have introduced a new property, bounded accumulated-error, and we have shown that, if all resource agents have this property, the performance of the overall grid-control system improves in several ways. Hence, we conclude that every resource agent should have c -bounded accumulated-error by design, for some appropriate (and scenario-specific) c .

We hope that this theoretically-oriented work will contribute to the development of a practical and scalable solution for controlling power grids with a significant fraction of renewable-energy sources.




Fig. 8. Plots of the accumulated error of the heater agent. In the absence of error diffusion (top figure), the accumulated error grows linearly with time. With error diffusion (bottom figure), the accumulated error is bounded from above and below.




Fig. 9. Time-averaged sequences of requested and implemented setpoints belonging to the heater agent. In the absence of error diffusion (top figure), the implemented setpoint does not converge to the objective, where the objective is the minimizer of the cost function. With error diffusion (bottom figure), the sequence of implemented setpoint converges towards the objective.

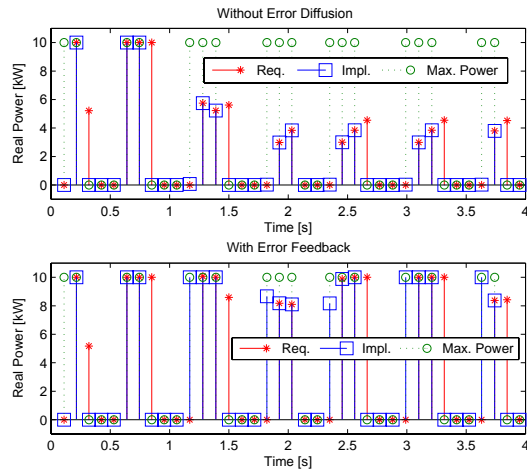


Fig. 10. Sequence of requested vs. implemented setpoints belonging to the PV agent. The bottom figure shows that error diffusion helps to maximise utilization of the PV.




Fig. 11. Plot of the accumulated error of the PV agent. Like in the heater-agent case, the accumulated error grows unboundedly in the absence of error diffusion (top figure), whereas the accumulated error is bounded with error diffusion.




Fig. 12. Time-averaged sequence of requested vs. implemented setpoints belonging to the PV agent. Also this plot shows that error diffusion helps to increase the utilization of the PV.

While the motivation for this work originated from problems that are specific to our application, we believe that our results from Section VII are general enough to be of independent interest.

REFERENCES

- [1] A. Bernstein, L. Reyes-Chamorro, J.-Y. Le Boudec, and M. Paolone, "A composable method for real-time control of active distribution networks with explicit power setpoints. Part I: Framework," *Electric Power Systems Research*, vol. 125, pp. 254–264, 2015.
- [2] L. Reyes-Chamorro, A. Bernstein, J.-Y. Le Boudec, and M. Paolone, "A composable method for real-time control of active distribution networks with explicit power setpoints. Part II: Implementation and validation," *Electric Power Systems Research*, vol. 125, pp. 265–280, 2015.
- [3] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grey scale," in *Proc. Dig. SID International Symp.*, Los Angeles, California, 1975, p. 3637.
- [4] D. Anastassiou, "Error diffusion coding for a/d conversion," *Circuits and Systems, IEEE Transactions on*, vol. 36, no. 9, pp. 1175–1186, Sep 1989.
- [5] R. Gray, "Oversampled sigma-delta modulation," *Communications, IEEE Transactions on*, vol. 35, no. 5, pp. 481–489, May 1987.

- [6] R. Adams and R. Schreier, "Stability theory for delta-sigma modulators," in *Delta-Sigma Data Converters: Theory, Design, and Simulation*, S. R. Norsworthy, R. Schreier, and G. C. Temes, Eds. New Jersey: Wiley, 1996.
- [7] G. T. Costanzo, A. Bernstein, L. E. Reyes Chamorro, H. W. Bindner, J.-Y. Le Boudec, and M. Paolone, "Electric Space Heating Scheduling for Real-time Explicit Power Control in Active Distribution Networks," in *IEEE PES ISGT Europe 2014*, 2014.
- [8] S. Papathanassiou, N. Hatzigiorgiou, and K. Strunz, "A benchmark low voltage microgrid network," in *Proceedings of the CIGRÉ Symposium "Power Systems with Dispersed Generation: technologies, impacts on development, operation and performances"*, Apr. 2005, Athens, Greece.